BAYESIAN STATISTICS 7, pp. 000–000
J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid,
D. Heckerman, A. F. M. Smith and M. West (Eds.)

A Hierarchical Model for Estimating the Reliability of Complex Systems

VALEN E. JOHNSON
Los Alamos National Laboratory, USA
valen@lanl.gov

TODD L. GRAVES
Los Alamos National Laboratory, USA
tgraves@lanl.gov

MICHAEL S. HAMADA
Los Alamos National Laboratory, USA
hamade@lanl.gov

C. SHANE REESE
Brigham Young University, USA
reese@lanl.gov

SUMMARY

We describe a hierarchical model for assessing the reliability of multi-component systems. Novel features of this model are the natural manner in which failure data collected at either the component or subcomponent level is aggregated into the posterior distribution, and pooling of failure information between similar components. Binary regression models are used to augment the model to account for the degradation of system performance with respect to time or other environmental factors. An example involving the performance of an anti-aircraft missile defense system illustrates the methodology.

Keywords: AGGREGATION ERROR, FAILURE MODEL, MULTI-COMPONENT SYSTEM, SYSTEM RELIABILITY.

1. BACKGROUND

This paper addresses the integration of component, subsystem and system data and prior expert opinion to assess system reliability as it changes over time. Two problems in reliability which separately have received much attention in the literature are thus combined: (1) the integration of available information at various levels to assess system reliability and (2) estimating reliability growth or degradation. Methodology for integrating available information in a consistent fashion has proven problematic, and this paper describes a Bayesian hierarchical model that resolves this difficulty. For simplicity, we restrict discussion to systems in which components or subcomponents may be regarded as either functional or not.

To provide context, it is useful to begin with a review of related research in Bayesian system reliability. Most relevant to the model considered here are the papers by Martz, Waller and Fickas (1988) and Martz and Waller (1990), where complex systems, comprised of series and parallel subcomponents, were modeled using beta priors and binomial likelihoods at component, subsystem and system levels. Within this framework, an "induced" higher-level prior was obtained by propagating lower-level posteriors up through the system fault diagram, and combining these posteriors with "native" higher-level priors to obtain an induced prior at the next system level. The induced priors were then approximated by beta distributions using a methods-of-moments type procedure. The combination of native priors and posterior distributions obtained from lower-level system data, both of which were expressed as beta distributions, was accomplished by expressing the induced priors as a beta distributions with parameters representing a weighted average of the constituent beta densities. This process was propagated through subsequent system levels to obtain an approximation to the joint posterior distribution on the total system reliability.

Many other reliability models are not able to account for prior expert opinion and data when such information is simultaneously obtained at several levels within a system. However, Springer and Thompson (1966, 1969), and Tang, Tang and Moskowitz (1994, 1997) have provided exact, and in complicated settings, approximate system reliability distributions based on binomial data by propagating component posteriors through the system's fault diagram. Thompson and Chang (1975), Chang and Thompson (1976), Lampkin and Winterbottom (1983) and Winterbottom (1994) employed approximations for system reliability distributions based on exponential lifetimes rather than binomial data. Others have proposed methods for evaluating or bounding moments of the system reliability posterior distribution (Cole (1975), Mastran (1976), Dostal and Iannuzzelli (1977), Mastran and Singpurwalla (1978), Barlow (1985), Natvig and Eide (1987), Soman and Misra (1993)); the first moment provides an estimate of system reliability. Moment estimators have also been used in the beta approximations employed by Martz, Waller and Fickas (1988) and Martz and Waller (1990). In a somewhat different approach, Soman and Misra (1993) proposed distributional approximations based on maximum entropy priors.

Numerous models have, of course, also been proposed for modeling non-binomial data. Thompson and Chang (1975), Chang and Thompson (1976), Mastran (1976), Mastran and Singpurwalla (1978), Lampkin and Winterbottom (1983), and Winterbottom (1994) considered models for exponential lifetime data, while Hulting and Robinson (1990, 1994) examined Weibull models. Poisson count data, where the number of units failing in a specified period, are discussed in Hulting and Robinson (1990), Sharma and Bhutani (1992), Hulting and Robinson (1994), Sharma and Bhutani (1994), and Martz and Baggerly (1997). Currit and Singpurwalla (1988) and Bergman and Ringi (1997a) considered dependence between components introduced through common operating environments. Bergman and Ringi (1997b) incorporated data from non-identical environments.

In many previously-defined reliability models, a logical difficulty arises when prior information and data are combined at distinct component levels. Bier (1994) discussed this difficulty, which arises because data integration may be accomplished in one of several ways. In one approach, component-level priors and data are propagated upward to higher system levels in order to obtain a system-level posterior. In another, component-level priors are propagated to the system level, where they are combined (only) with system-level data to obtain a posterior distribution on the system reliability. Unfortunately, these approaches generally yield different posterior distributions on the system reliability. This effect is known as aggregation error.

Johnson, Graves, Hamada and Reese (2001) proposed a Bayesian hierarchical model which bypasses the consistency problems addressed by Bier, but yet is capable of integrating available information at all component levels to yield a posterior distribution on system reliability. In this article, we extend that approach to account for system reliability degradation over time.

Previous degradation models for system reliability have typically restricted attention to settings in which only system-level data are available (e.g., Fries and Sen (1996), Nolander and Dietrich (1994), and Sohn (1996)). An exception to this trend is Robinson and Dietrich (1988), who modeled component-level data collected during system development using exponential lifetime assumptions and decreasing failure rates. Our approach differs from that taken by Robinson and Dietrich in that we utilize binary regression models obtained at multiple component levels to model aging processes.

An outline for the remainder of this article is as follows. In Section 2, we review the baseline reliability model described in Johnson, Graves, Hamada and Reese (2001). This model is illustrated with an application to anti-aircraft missile system data in Section 3. The extension of the model to account for time degradation (or other component-level covariates) is described in Section 4. We conclude with a summary of results and suggestions for future work in Section 5.

2. METHODOLOGY

To illustrate the baseline model, consider Figure 1, which depicts a fault tree for an anti-aircraft missile system. The general features illustrated in this figure include the composition of a system by multiple subsystems, and the composition of these subsystems by further subsystems and components. In general, we assume that binomial data and prior expert opinion are available at different system levels, and that our primary goal in modeling such systems is the evaluation of the probability that a system (missile) drawn at random from the stockpile functions. Secondary goals might involve advising stockpile/inventory managers of the utility of conducting full-system or component-level tests to evaluate this probability, and to identify subsystems for which additional data might best be collected to improve estimation of overall system reliability.

Several sources of information relevant to estimating system reliability are considered. The first is binomial data collected from actual component or subsystem tests. In the augmented degradation model, the age of the component at the time of the tests or other covariate information is also assumed to be available. The second source of information takes the form of expert opinion regarding the probability that a specific component or subsystem fails. This information is accompanied by relevant covariate values in the augmented degradation model. A third, less precise source of information is expert opinion regarding the similarity of the failure probabilities of groups of components within the system or across different systems. For example, in the missile system depicted above, an expert may assert that the reliability of the missile battery is similar to the reliability of a battery in a related missile system, or that reliabilities of the eject and flight motors are similar. However, the expert may not have knowledge regarding the specific probability that any component within a group of similar components functions. Finally, we incorporate the statistical notion that terminal nodes (i.e., components in the fault tree having no subcomponents themselves) may also be grouped into sets of comparably reliable components without the guidance of actual

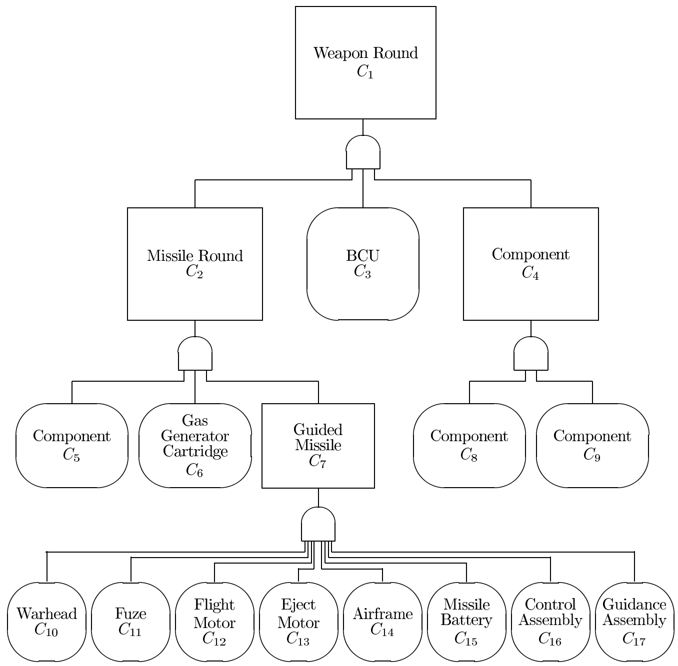

Figure 1. *Reliability Fault Tree Diagram for an Anti-Aircraft Missile System*

expert opinion. In the baseline model, such information is modeled via an exchangeability assumption on the terminal probabilities themselves, while in the degradation model this assumption takes the form of an exchangeability assumption on binomial regression coefficients.

To model these sources of information, we first assume that the failure probabilities of components in distinct branches of the fault tree are conditionally independent, and that the success of the system requires successful functioning of all components. Extensions to systems that include redundant components, or in which component failures are not independent, are discussed in the summary. Nodes in the reliability diagram are labeled C_i, where i indicates the component or subcomponent index. The function $a(i)$ provides the parent component (or system) containing (sub)component i, while $g(i,m)$ indicates the group of components that expert m asserts have similar failure rates to component i. We let p_i denote the probability that component C_i functions when the missile is fired. The set of components for which test data is available is denoted by S_0, and within this set x_i denotes the number of times component i functioned successfully in n_i trials. In the baseline model, aging effects are not considered, making a simple binomial likelihood appropriate for modeling (x_i, n_i).

In many actual applications, expert opinion plays a potentially important role in assessing system reliability, particularly in large complex systems for which data collected on individual subcomponents may be sparse. Furthermore, expert opinion may be available from several experts, and the quality of information obtained from each expert may vary. In the baseline model, we therefore assume that the prior density obtained from expert m concerning a specific value of p_i takes the form of a beta density, and we let the set of combinations of (i,m) for which expert opinion is available be denoted by S_1. More specifically, we assume that the net contribution in the joint

Table 1. *Notation used in model definition*

C_i	Component i in system fault diagram
p_i	Probability that component i functions
$a(i)$	Parent of component i
$g(i,m)$	The group to which expert m assigns component i
S_0	Set of components for which test data is available
x_i	Number of successful trials of component i
n_i	Total number of trials of component i
p_{ij}	Success probability of component i under conditions z_{ij}
x_{ij}	Number of successful trials of component i under z_{ij}
n_{ij}	Total trials of component i under conditions z_{ij}
z_{ij}	Covariate vector
$\boldsymbol{\beta}_i$	Regression coefficient for ith terminal node
S_1	Components for which specific expert opinion is available
$\pi_{i,m}$	Expert m's point estimate of p_i
N_m	Beta parameter describing expert m's precision
$\alpha_m,\ \beta_m$	Gamma distribution parameters in prior on N_m
S_2	Components for which grouping, information is available
$\rho_{m,g}$	Central value of beta density on success probabilities for components in group g defined by expert m
K_m	Beta dispersion parameter for component probabilities in group g around $\rho_{m,g}$
$\delta_{g,m},\ \epsilon_{g,m}$	Beta hyperparameters in prior for $\rho_{m,g}$
$\zeta_m,\ \eta_m$	Gamma hyperparameters in prior for K_m
S_3	Set of terminal nodes
ϱ_0	Central value of beta density assumed for terminal nodes
J_0	Pior beta dispersion parameter for ϱ_0
$\tau_0,\ \phi_0$	Gamma hyperparameters in prior for K_m
$\psi_0,\ \omega_0$	Beta hyperparameters in prior on ϱ_0

posterior density arising from such prior information is

$$\frac{\Gamma(N_m+2)}{\Gamma(N_m\pi_{i,m}+1)\Gamma[N_m(1-\pi_{i,m})+1]p_i^{N_m\pi_{i,m}}(1-p_i)^{N_m(1-\pi_{i,m})}} \\ \equiv \mathrm{B}(p_i \,|\, N_m\pi_{i,m}+1, N_m(1-\pi_{i,m})+1). \tag{1}$$

In (1), $\pi_{i,m}$ represents expert m's point estimate of p_i, and N_m represents the precision of expert m. For concreteness, we assume that each expert precision parameter N_m is drawn from a gamma density with known parameters α_m and β_m, parameterized here as

$$\mathrm{G}(N_m \,|\, \alpha_m, \beta_m) = \frac{\beta_m^{\alpha_m}}{\Gamma(\alpha_m)} N_m^{\alpha_m-1}\exp(-\beta_m N_m).$$

Note that expert opinion is assumed to take the form of a binomial likelihood with a maximum at $\pi_{i,m}$ – this convention eliminates the possibility that the joint density specified on all model parameters is improper, and also implicitly handles the aggregation problem identified by Bier (1994) by simply treating expert opinion as data.

When prior information regarding component success probabilities is unknown, but expert groupings of components are available, (1) is augmented in the baseline model by assuming that $\pi_{i,m}$ is replaced by $\rho_{m,g}$, where $\rho_{m,g}$ represents the common, but unknown, success probability assigned by expert m to components in group g (i.e., components for which $g(i,m) = g$). The contribution to the joint posterior distribution on model parameters from such information is assumed to take the form

$$\prod_{(i,m)\in S_2} \mathrm{B}(p_i\,;\, K_m\rho_{m,g} + 1, K_m(1-\rho_{m,g}) + 1). \tag{2}$$

Here, S_2 denotes the combinations of (i,m) for which such grouping information is available.

As in (1), the parameter K_m is assumed to be drawn *a priori* from a gamma density having parameters ζ_m and η_m. The prior success parameter $\rho_{m,g}$ is assumed to be drawn from a beta density with known parameters $\delta_{g,m}$ and $\epsilon_{g,m}$, respectively.

Finally, for terminal nodes in the fault tree a hierarchical prior specification may be obtained by further assuming that each terminal node's success probability is drawn from a beta density with parameters $J_0\varrho_0$ and $J_0(1-\varrho_0)$. The set of terminal nodes is denoted by S_3.

For notational simplicity, we assume that all terminal nodes are, *a priori*, exchangeable, but this restriction may be relaxed by using expert judgment to group the terminal nodes in a manner similar to that used in the specification of (2). In that case, J_0 and ϱ_0 would be subscripted with the appropriate prior group. The parameter J_0 is assumed drawn from a gamma density with parameters τ_0 and ϕ_0; ϱ_0 is assumed *a priori* to be drawn from a beta density with parameters ψ_0 and ω_0.

As discussed in the previous section, combining data and prior information at different levels within a reliability diagram has often proven problematic, both from the perspectives of computational tractability and model consistency. Our solution to this conundrum is to simply re-express non-terminal node probabilities in terms of terminal node probabilities using deterministic relations derived from an examination of the system reliability diagram. For example, from Figure 1, it is evident that the probability that the guided missile component functions, p_7, is equal to the product of the probabilities that the warhead (p_{10}), fuze (p_{11}), flight motor (p_{12}), eject motor (p_{13}), airframe (p_{14}), missile battery (p_{15}), control assembly (p_{16}), and guidance assembly(p_{17}) all function. Thus,

$$p_7 = \prod_{i=10}^{17} p_i \tag{3}$$

and, for example, the prior specification on p_7 is interpreted as a prior specification on this product:

$$\begin{aligned} f_{7,m}(p_7 \mid \pi_{7,m}, K_m) &\equiv f_{7,m}\Big(\prod_{i=10}^{17} p_i \mid \pi_{7,m}, N_m\Big) \\ &\propto \left[\prod_{i=10}^{17} p_i\right]^{N_m\pi_{7,m}} \left[1 - \prod_{i=10}^{17} p_i\right]^{N_m(1-\pi_{7,m})} . \end{aligned}$$

Note that variable substitutions based on the reliability diagram do not uniquely identify a joint distribution on the terminal node probabilities, in this case p_{10} through p_{17}. However, together with the assumption that the distributions of these probabilities are defined with respect to Lebesgue measure on the unit interval and the given hierarchical specification, such substitutions do yield a uniquely defined joint distribution on these parameters.

Combining these assumptions leads to a joint posterior distribution on the baseline model parameters proportional to

$$
\begin{aligned}
&[\boldsymbol{p}, N, \boldsymbol{\rho}, K, \boldsymbol{\varrho}, J \,|\, x, n, \boldsymbol{\pi}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\zeta}, \boldsymbol{\eta}, \boldsymbol{\delta}, \boldsymbol{\epsilon}, \boldsymbol{\tau}, \boldsymbol{\phi}, \boldsymbol{\psi}, \boldsymbol{\omega}] \propto \\
&\quad \times \prod_{i \in S_0} p_i^{x_i}(1-p_i)^{n_i} \times \prod_{m:\exists(i,m)\in S_2} \mathrm{G}(K_m\,;\,\zeta_m, \eta_m) \\
&\quad \times \prod_{(i,m)\in S_1} \mathrm{B}(p_i\,;\,N_m\pi_{i,m}+1, N_m(1-\pi_{i,m})+1) && (4) \\
&\quad \times \prod_{(i,m)\in S_2} \mathrm{B}(p_i\,;\,K_m\rho_{m,g}+1, K_m(1-\rho_{m,g})+1) && (5) \\
&\quad \times \prod_{i\in S_3} \mathrm{B}(p_i\,;\,J_0\varrho_0, J_0(1-\varrho_0)+1) && (6) \\
&\quad \times \prod_{m:\exists(i,m)\in S_2} \mathrm{B}(\rho_{m,g}\,;\,\delta_m, \epsilon_m) \times \mathrm{B}(\varrho_0\,;\,\psi_0, \omega_0) \\
&\quad \times \mathrm{G}(J_0\,;\,\tau_0, \psi_0) \times \prod_{m:\exists(i,m)\in S_1} \mathrm{G}(N_m\,;\,\alpha_m, \beta_m). && (7)
\end{aligned}
$$

In this expression, values of non-terminal node probabilities are assumed to be expressed in terms of the appropriate functions of terminal node probabilities, as defined from the system fault diagram.

An examination of the contributions to the joint posterior distribution arising from the three types of prior information (4–6) reveals obvious similarities, but there are also important distinctions between these parameterizations. For example, in (4), the value of N_m represents the precision of the expert's opinion, while in (5) and (6), K_m and J_0 describe the similarity of item reliabilities within a grouping.

2.1 *Hierarchical prior model*

The hierarchical prior model on the terminal node probabilities plays a crucial role in rendering estimates of the overall system reliability insensitive to the level of detail included in the system fault diagram. As an illustration of this point, consider a simple system comprised of three components, and suppose that a single binomial observation with 4 successes and 1 failure is observed at the system level. Then without a hierarchical specification on the component probabilities and under the model assumptions stated above, the posterior distribution on the system reliability would be proportional to where the system reliability, p_1, is assumed to equal $p_2p_3p_4$.

With the implied uniform distribution on p_2–p_4, the posterior mean of p_1 in this model is 0.507; when the system is not decomposed into subsystems and a uniform prior is assumed on p_1, the posterior mean on p_1 (with a uniform prior) is .714. Furthermore, under such naive model specifications, the bias attributable to adding subcomponents to the fault tree becomes more severe as the number of subcomponents in the system increases.

In contrast, the hierarchical prior specification on p_2–p_4 with $\psi_0 = \omega_0 = 0.5$ results in a posterior mean of 0.718 for p_1, while the same specification with $\psi_0 = \omega_0 = 1.0$ results in a posterior mean of 0.687. Both estimates are largely insensitive to the number of subcomponents specified for the system.

2.2 *Estimation strategies for the baseline model*

The joint distribution on model parameters specified in (7) does not lend itself to analytical evaluation of the system or component reliabilities. However, a component-wise Metropolis-Hastings algorithm can be implemented in a relatively straightforward way. In our version of such a scheme, we used a random-walk Metropolis-Hastings algorithm with Gaussian proposal densities specified on the logistic scale for the terminal node probabilities, as well as for ϱ_0 and $\rho_{m,g}$. Precision parameters were similarly updated using a random-walk Metropolis-Hastings scheme with Gaussian increments specified on the logarithmic scale. The resulting Metropolis-Hastings algorithm was implemented using a general-purpose Java MCMC system developed at Los Alamos National Laboratory (Graves, 2001).

3. ANALYSIS OF ANTI-AIRCRAFT MISSILE DATA

Anti-aircraft missiles are intended to provide defense from attacking enemy aircraft. Currently, the United States has over 15 different anti-aircraft missiles in its arsenal.

Each of these weapons is comprised of numerous components and subsystems, many of which are depicted schematically for a selected weapon system in Figure 1. Data available for assessing the reliability of this particular system include 45 observations on each of the components $C_4, C_5, C_6, C_{11}, C_{12}, C_{13}, C_{15}$ and C_{16}, and 126 tests of component C_3. The bulk of the test data, however, was performed at the system level, where over 1400 tests were performed. This situation is atypical of most applications in which component-level tests dominate, but this feature of the data offers an ideal opportunity for us to test our reliability model by comparing results obtained both with and without the full-system data. No component-level tests were performed on components $C_2, C_7, C_8, C_9, C_{10}, C_{14}$, and C_{17}.

Upon conferring with an expert, reliability classes were formed as follows. Components 2-4 were assigned to Group 1, Components 5-9 to Group 2, Components 10-17 to Group 3, and Components 5, 6, and 8-17 to Group 4. Informative priors with common, fixed group means (π_1-π_3) and a single, common precision parameter ($N_{1,2,3}$) were assumed for each of Groups 1-3. A common precision parameter was incorporated for each group owing to the fact that a single expert provided all prior information. A hierarchical prior with unknown mean and precision parameter (ϱ_0 and J_0, respectively) was assumed for components in Group 4. Also, gamma priors with parameters (5,1) were posited for both precision parameters ($N_{1,2,3}$ and J_0), and a non-informative prior ($\psi_0 = \omega_0 = 0.5$) was assumed for ϱ_0.

Applying the model discussed in Section 2, we obtained the posterior distributions on the component reliabilities for each of the components and the expert precision parameters. The system reliability posterior distributions with the system data included and system data excluded are plotted in Figure 2. We note the agreement between the two posterior distributions (full system tests included vs. full system tests excluded). In every case, the 95% posterior probability region calculated by excluding full system data includes the posterior distribution obtained with the full system data. The scales from these plots have been removed due to classification concerns, although it can be noted that each plot contains a subinterval of $(0, 1)$ of length 0.1 or less. Also of interest is the posterior distribution for the expert precision parameter $N_{1,2,3}$. The posterior mean for this distribution is 12.2. This suggests that the expert's opinion is worth approximately 12 full system tests. Given the prior mean of 5, we conclude that the expert was reasonably well calibrated with the system structure and data.

3.1 *Diagnostics*

Two concerns commonly encountered in modeling system-level reliabilities using fault diagrams like that depicted in Figure 1 involves the extent to which different components function independently and whether system (or subsystem) reliability decreases when components are assembled. A simple cross-validation diagnostic useful for assessing the importance of these influences can be constructed by iteratively omitting data collected at each node from the estimation procedure, and then examining the predictive distribution for the omitted datum.

Such a procedure was applied to data obtained for this missile system, and resulted in an estimate of 0.83 for the predictive probability of observing fewer successes at the system-level than were actually observed. It therefore seems that there is little evidence to support the notion that the reliability of the system was degraded as components were assembled and required to operate as a unit.

There was, however, some indication of model lack-of-fit at the subcomponent level. For components 10 and 17, the predictive probability for observing fewer successes than were obtained at these nodes was approximately 0.035. The same number of failures was observed at each of these components, and these two components had the highest failure rate of any components in the system. Model lack-of-fit in this instance might thus be attributed to the fact that the hierarchical mean estimated for the terminal nodes, ϱ_0, increased substantially when the datum for either of these nodes was omitted, resulting in an overly optimistic estimate of this probability. Possible remedies for such model inadequacy would be to stochastically decrease the prior assigned to the value of J_0, or to introduce a separate hierarchical group for these nodes. In this case, neither remedy appeared to substantially affect estimates of system reliability in follow-on sensitivity analyses.

4. EXTENSION TO DEGRADATION MODELS

In many complex systems, reliabilities of system components degrade with age, and such degradation processes can be modeled naturally within the hierarchical framework described above. For purposes of illustration, we describe this extension within the context of logistic regression models on component reliabilities; generalizations to broader classes of regression models follow along similar lines.

In the baseline model, all component reliabilities are specified in terms of the reliabilities of terminal nodes, thus making it possible to consistently incorporate information collected at multiple component levels. A similar approach is adopted for modeling the degradation of component reliabilities over time. Specifically, for each terminal node in the system, say node i, we assume that

$$\log\left(\frac{p_{ij}}{1-p_{ij}}\right) = z'_{ij}\boldsymbol{\beta}_i,$$

where z_{ij} denotes a known vector of covariates relevant for predicting p_{ij} (the success probability of component i under conditions z_{ij}) and $\boldsymbol{\beta}_i$ denotes a regression coefficient specific to terminal node i. In the case of the missile data discussed later, the vectors z contain a constant term (intercept) and component age at time of testing or prior specification.

The hierarchical structure assumed for the component reliabilities in the baseline model is extended to the regression setting by assuming that, for terminal nodes within

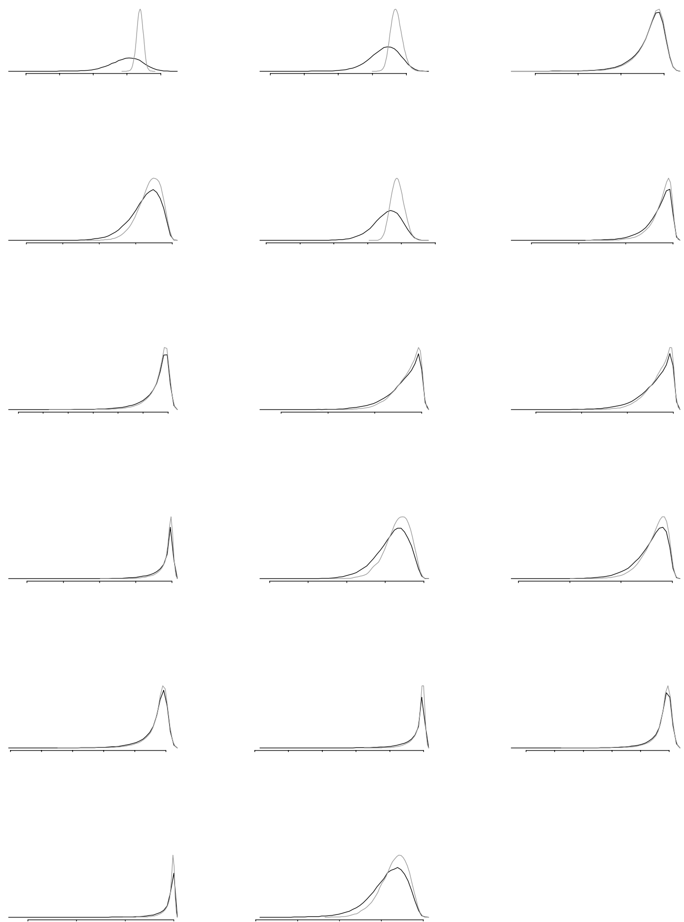

Figure 2. *Posterior distributions for the reliability of the system represented in Figure 1. In each pair of plots, the more peaked curve represents the marginal posterior density based on all test data, while the more dispersed curves represent the marginal posterior density using only component-level data (i.e., excluding system-level tests).*

a common grouping, the vectors β_i are drawn from a multivariate normal distribution with mean, say, α and covariance matrix C. In this application, vague priors are assumed for α and C.

Specification of prior expert opinion is also modified to account for changes in component reliabilities over time. This is accomplished by substituting

$$\frac{\exp(z_{ij}'\boldsymbol{\beta})}{1+\exp(z_{ij}'\boldsymbol{\beta})}$$

for p_{ij} in equations (1) and (2) for an assumed value of the covariate vector z_{ij}. Often, z_{ij} is chosen to correspond to the state of a component at time 0.

Estimation of model parameters proceeds as in the baseline model, except that a random-walk Metropolis-Hastings update for the values of $\boldsymbol{\beta}_i$ replaces the corresponding updates of the values of p_{ij}.

4.1 *Application to missile data*

The success/failure data described in Section 3 were accompanied by the age, in months, of each subsystem at the time of the tests. Tests were conducted at 19 distinct times for components $C_4, C_5, C_6, C_{11}, C_{12}, C_{13}, C_{15}$ and C_{16}, while the 126 tests of Component 3 were conducted at 23 distinct times. The 1,400 system-level tests were performed at 106 distinct system ages. Ages at which system tests were conducted ranged from 0 to 143 months, though most system-level data was collected from systems less than 3 years old.

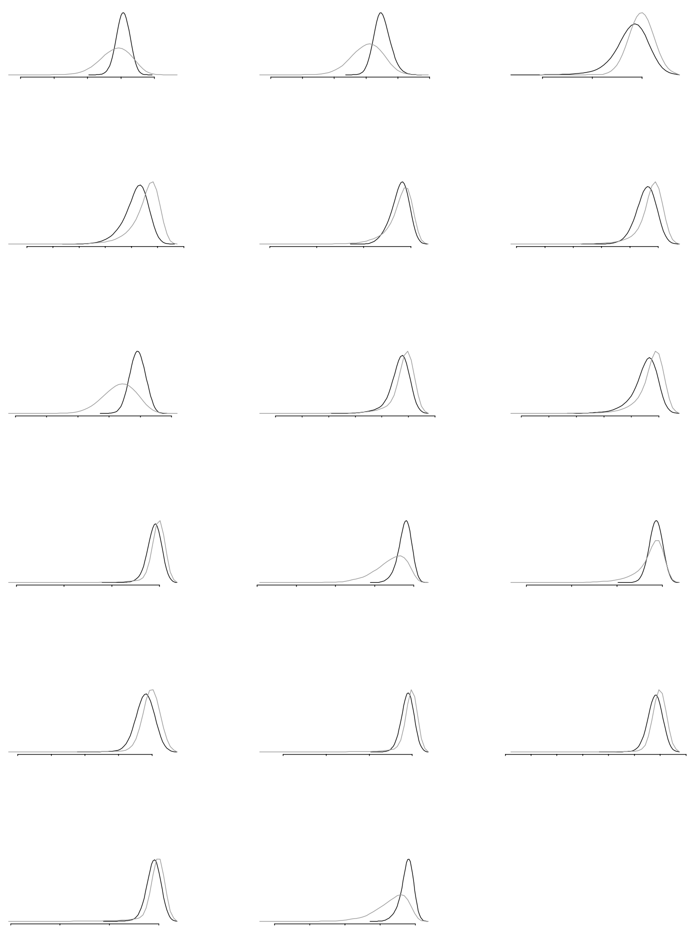

Figure 3. *Posterior distributions for the reliability of the missile-system components at time 0 and 120 months. In each case, the more disperse density corresponds to the posterior density estimated for 120 months.*

A plot of the marginal posterior densities on the reliabilities of components at different levels within the system is depicted in Figure 3. These plots were obtained by assuming that all prior opinion elicited in the baseline model applied at time 0 and using all system- and component-level data. As expected, greater uncertainty is associated with the reliability of most estimates at 10 years, owing to the comparatively high posterior uncertainty in many of the values of the regression coefficient corresponding to system age ($\boldsymbol{\beta}_i$ in (8)). Note also that the posterior mean of the reliability of most components decreases gradually over time, again as is expected.

Plots of the marginal posterior distributions of the slope parameters of the terminal nodes are depicted in Figure 4. These plots highlight the extent to which many of the individual regression parameters are only weakly identifiable. In particular, for terminal nodes and subcomponents for which data were collected at a limited number of time points, and for which no failures were observed, many logistic regression curves provide nearly equivalent fits in the vicinity of the observed data. This phenomenon is further exacerbated at the subsystem level when data at lower-level components is sparse.

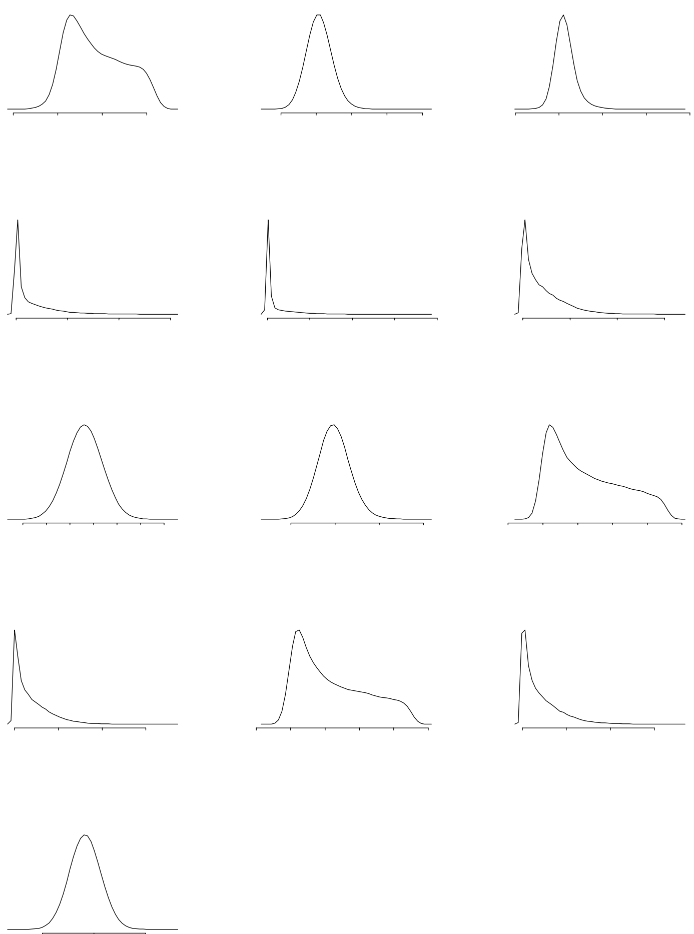

Figure 4. *Posterior distributions for the logistic slope parameters for terminal nodes. Note that Components 3, 13, and 15 experienced no failures in component-level testing.*

5. CONCLUSIONS

The proposed hierarchical model offers several advantages over existing models for system reliabilities. Among these are the ease of including diverse sources of information at different levels of the system into the model for overall system reliabilities, a coherent framework for incorporating multiple sources of prior expert opinion through the treatment of expert opinion as (imprecisely-observed) data, and the natural elimination of aggregation errors through the definition of non-terminal probabilities using the assumed structure of the system fault tree and terminal node probabilities.

A simplistic form of our hierarchical model for reliability was described in this paper. In future work we plan to extend this framework to include non-serial systems and extensions of the model to account for dependencies between component failures within a system or subsytem.

REFERENCES

Barlow, R.E. (1985) "Combining Component and System Information in System Reliability Calculation," in *Probabilistic Methods in the Mechanics of Solids and Structures*, Eds. S. Eggwertz and N.L. Lind, Springer-Verlag, Berlin, 375–383.

Bergman, B. and Ringi, M. (1997a), "Bayesian System Reliability Prediction," *Scandinavian J. Statist.* **24**, 137–143.

Bergman, B. and Ringi, M. (1997b), "System Reliability Prediction Using Data from Non-identical Environments," *Reliability Engineering and System Safety*, **58**, 183–190.

Berkson, J. (1944), "Application of the Logistic Function to Bio-assay," *J. Amer. Statist. Assoc.* **39**, 357–365.

Bier, V.M. (1994), "On the Concept of Perfect Aggregation in Bayesian Estimation," *Reliability Engineering and System Safety*, **46**, 271–281.

Chang, E.Y. and Thompson, W.E. (1976), "Bayes Analysis of Reliability of Complex Systems," *Operations Research* **24**, 156–168.

Cole, P.V.Z. (1975), "A Bayesian Reliability Assessment of Complex Systems for Binomial Sampling," *IEEE Trans. Reliability*, **R-24**, 114–117.

Currit, A. and Singpurwalla, N.D. (1988), "On the Reliability Function of a System of Components Sharing a Common Environment," *J. Appl. Probability* **26**, 763–771.

Dostal, R.G. and Iannuzzelli, L.M. (1977), "Confidence Limits for System Reliability When Testing Takes Place at the Component Level," in *The Theory and Applications of Reliability* (Vol. 2), Academic Press, New York, 531–552.

Fries, A. and Sen, A. (1996), "A Survey of Discrete Reliability-Growth Models," *IEEE Trans. Reliability*, **45**, 582–604.

Graves, T.L. (2001), "YADAS: An Object-Oriented Framework for Data Analysis Using Markov Chain Monte Carlo," Los Alamos National Laboratory Technical Report, LA-UR-01-4804.

Hulting, F.L. and Robinson, J.A. (1990), "A Bayesian Approach to System Reliability," Research Publication No. GMR-7110, General Motors Research Laboratories, Warren, MI.

Hulting, F.L. and Robinson, J.A. (1994), "The Reliability of a Series System of Repairable Subsystems: A Bayesian Approach," *Naval Research Logistics*, **41**, 483–506.

Johnson, V.E., Graves, T.L., Hamada, M. and Reese, C.S. (2001), "A Hierarchical Model for Estimating the Reliability of Complex Systems," Los Alamos National Lab Technical Report LA-UR-01-6915.

Lampkin, H. and Winterbottom, A. (1983), "Approximate Bayesian Intervals for the Reliability of Series Systems from Mixed Subsystem Test Data,", *Naval Research Logistics*, **30**, 313–317.

Martz, H.F. and Baggerly, K.A. (1997), "Bayesian Reliability Analysis of High-Reliability Systems of Binomial and Poisson Subsystems," *International J. of Reliability, Quality and Safety Engineering*, **4**, 283–307.

Martz, H.F. and Waller, R.A. (1990), "Bayesian Reliability Analysis of Complex Series/Parallel Systems of Binomial Subsystems and Components," *Technometrics*, 32, 407–416.

Martz, H.F., Waller, R.A. and Fickas, E.T. (1988), "Bayesian Reliability Analysis of Series Systems of Binomial Subsystems and Components," *Technometrics* **30**, 143–154.

Mastran, D.V. (1976), "Incorporating Component and System Test Data Into the Same Assessment: A Bayesian Approach," *Operations Research* **24**, 491–499.

Mastran, D.V. and Singpurwalla, N.D. (1978), "A Bayesian Estimation of the Reliability of Coherent Structures," *Operations Research* **26**, 663–672.

Natvig, B. and Eide, H. (1987), "Bayesian Estimation of System Reliability," *Scandinavian J. Statist.* **14**, 319–327.

Nolander, J.L. and Dietrich, D.L. (1994), "Attribute Data Reliability Decay Models," *Microelectronics Reliability*, **34**, 1565–1596.

Robinson, D. and Dietrich, D. (1988), "A System-Level Reliability-Growth Model," *Proc. Annual Reliability and Maintainability Symposium*, 243–247.

Sharma, K.K. and Bhutani, R.K. (1994), "Bayesian Reliability Analysis of a Parallel System," *Microelectronics Reliability*, **34**, 761–763.

Sharma, K.K. and Bhutani, R.K. (1992), "Bayesian Reliability Analysis of a Series System," *Reliability Engineering and System Safety*, **53**, 227–230.

Soman, K.P. and Misra, K.B. (1993), "On Bayesian Estimation of System Reliability," *Microelectronics Reliability*, **33**, 1455–1459.

Sohn, S.Y. (1996), "Growth Curve Analysis Applied to Ammunition Deterioration," *J. Qual. Technol.* **28**, 71–80.

Springer, M.D. and Thompson, W.E. (1966), "Bayesian Confidence Limits for the Product of N Binomial Parameters," *Biometrika* **53**, 611–613.

Springer, M.D. and Thompson, W.E. (1969), "Bayesian Confidence Limits for System Reliability," in *Proc. 1969 Annual Reliability and Maintainability Symposium*, Institute of Electrical and Electronics Engineers, New York, 515–523.

Tang, J., Tang, K. and Moskowitz, H. (1994), "Bayes Credibility Intervals for Reliability of Series Systems with Very Reliable Components," *IEEE Trans. Reliability*, **43**, 132–137.

Tang, J., Tang, K. and Moskowitz, H. (1997), "Exact Bayesian Estimation of System Reliability from Component Test Data," *Naval Research Logistics*, **44**, 127–146.

Thompson, W.E. and Chang, E.Y. (1975), "Bayes Confidence Limits for Reliability of Redundant Systems," *Technometrics* **17**, 89–93.

Winterbottom, A. (1984), "The Interval Estimation of System Reliability from Component Test Data," *Operations Research* **32**, 628-640.